

Control Number : GOOG-AT-MDL-000881089

All Custodians : Ali Nasiri Amini, Bashar
Kachachi, Bonnie Pericolosi, Brent
Besson, Nirmal Jayaram, Vivek Rao

Custodian : Brent Besson

Date/Time Created : 1/24/2020 6:54 PM

Date/Time Sent :

MODIFICATION DATE : 9/14/2022 4:03 PM

File Name : DBM optimizations deep
dive_1wQvdhd0zu4A9zW0hUuIyz
Fd3uS08iMBGBdJiHmb9l-w.pptx



DV360 optimizations ENG deep dive

Confidential + Proprietary

1


DV360 five years ago

- Mostly fixed CPM manual bidding
 - Just after skyray
 - Auto-bidding products weren't effective or reliable
 - Auto-bidding adoption was ~15% and shrinking
- Exchanges were starting to move to non-second-price (including first-price) auctions
 - DV360 did not account for exchange auction mechanisms
 - 3PE accounted for over 50% of DV360 spend

We leveraged optimizations to revamp the product

- We completely revamped the auto-bidding offerings
 - Improved the underlying core models and technology to make the product more effective and reliable
 - Increased the suite of auto-bidding products available to advertisers
- Optimize DV360 buying on exchanges with adversarial auction mechanisms
 - This applies on both auto-bidding and fixed-bidding adgroups
 - From that perspective, 100% of DV360 has optimized bidding and is not just a “dumb pipe”

Impact on auto-bidding adoption



We have a suite of auto-bidding products

- Maximize conversions is by far the largest product
- Click and conversion based products account for 82% of auto-bidding spend

Product	Share
Maximize Conversions	54.3%
Maximize Clicks	17.5%
Maximize Active Views	9.9%
Target CPA	6.4%
Maximize AVOC	5.1%
Target CPC	4.2%
Custom Bidding	2.1%
Maximize TOS10	0.7%

Core improvements

Bidding optimization

- We first moved the underlying bidding technology to HDMI (to be covered in a separate deepdive)
- The bidding policy is the solution to an optimization problem to maximize advertiser value subject to constraints e.g.
 - Maximize Conversions
 - subject to
 - $\text{Spend} \leq \text{Budget}$
 - $\text{CPA} = \text{tCPA}$
- HDMI works by building a model of the market, while recognizing that we only see a partial view of the market, and leveraging our signals (e.g., pCvR) to identify the optimal bidding policy
- This launched doubled the conversions at the same tCPA

Bidding optimization

- ◆ While the technology was previously used in Adwords, several DV360 nuances needed to be handled
 - Advertisers pay per impression, so impression-level data needed to be used
 - A significant portion of DV360 spends is on 3PE, which introduces ecosystem-based challenges
 - A large portion of advertisers value CTCs and VTCs
 - Data sharing policy posed challenges - had to build siloed models using PEAR army (xbao@'s team) - separate deep-dive for predictions
- ◆ Recently, post AdX 1p migration, we have been able to get a full picture of the market because AdX shares HOBs, that has enabled us to make HDMI even more powerful


Budget handling

- Used to be a major pain point for many advertisers
- Earlier, budgets and bidding were treated as separate components, though they are “2 sides of the same coin”

Maximize Conversions
subject to
 $\text{Spend} \leq \text{Budget}$

- We embarked on a series of projects that rolled budget handling into the bid optimization framework, allowing us to significantly improve budget attainment and performance.

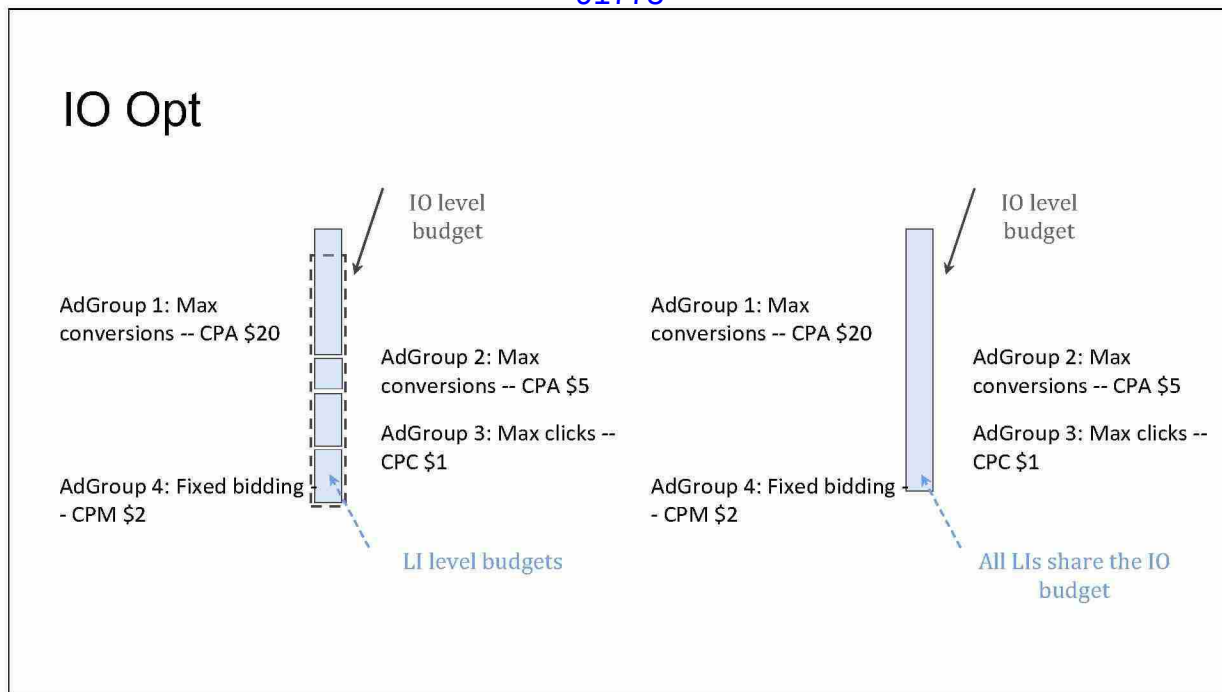
Bid lowering



IO opt


- DV360 allows advertisers to set budgets at the IO level as well as at the line item level
- Advertisers can set these LI level budgets manually or using Auto-Budget-Allocation (ABA)
- ABA was originally invented to help advertisers in a world where bidding wasn't optimized, but now DV360 is 100% optimized bidding...





IO opt adoption

- The budget allocation is a natural outcome of the solution to the optimization problem
- The optimization increased spend by 20% and CPD by 25%!
- Since launch in end of Q1 2019, IO opt has seen steep growth and is now at **\$1.2B ARR**



Key new products

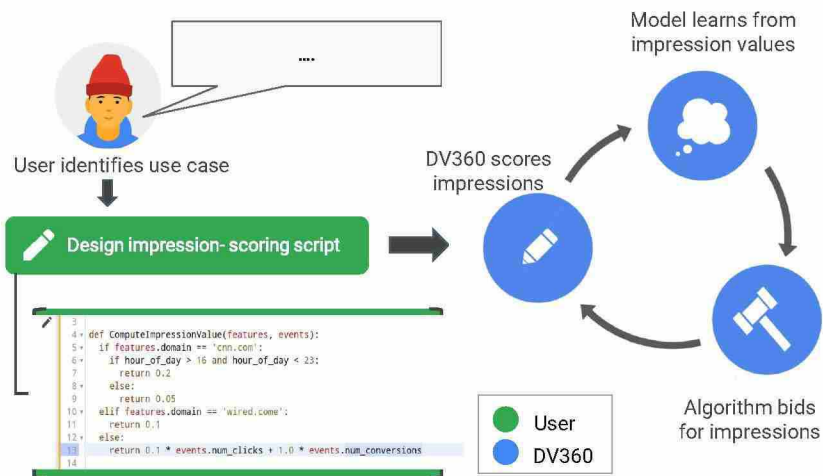
Brand/video optimizations

- The underlying bidding technology can be used in conjunction with several optimization goals.
- As a result, we have been able to introduce multiple new brand/video optimization goals that have had successful adoption
 - AVOC: Audible & Viewable on Completion
 - TOS10: Time on screen \geq 10 seconds
 - Active Views

Custom Bidding

- ◆ Agencies like features that help them show their value-add to their advertisers
- ◆ But providing them knobs that could come in the way of optimization and performance gains is detrimental.
- ◆ Custom bidding is one example of a feature where agencies and Google are aligned
 - Agencies have some information about what's valuable to their advertisers
 - Sometimes this is complex, and not easily expressible as clicks and conversions
 - In such cases, we offload the value attribution component to agencies, and use Google optimization technology to maximize that value.

Custom bidding



Custom Bidding

Trial results to date

- % of trials where test outperformed control: **73%**
- Aggregate performance uplift (averaged across trials): **+180%**
- 83 customers, current ARR \$60M

Use case	% of Trials	Avg. Perf. Lift**
Weighted conversions	49%	+205%
Transaction value	19%	+89%
Other*	32%	+207%

* Other includes use cases e.g. regression against offline metric, custom attention score, video completion on preferred domains, transaction value + weighted conversions, etc.

** Averaged across trials (not spend-weighted)

Pay-Per-Outcome (PPO)

- Two years ago, there was a goal to get performance parity between tCPA buying in DV360 and GDA
- There was no way to achieve this because of the differences in the buying modes
 - The Bernanke optimization framework in GDA leverages the arbitrage opportunities to win some impressions at high margins and therefore some at low margins, providing a significant advertiser value boost


Maximize Conversions
subject to
Spend \leq Budget
Per impression Margin = 15%

Ecosystem optimizations

1st price bidding

- Even before AdX moved to 1p, several exchanges had started deviating from second price auctions.
- Since DV360 fixed CPM did not have any optimizations at that time, the spend on 3PE increased disproportionately and came at the expense of advertiser performance
- Project Poirot was conceived to lower advertiser bids in first-price auctions to maximize advertiser surplus. Changes were made both to fixed CPM and auto-bidding advertiser bids.
- DV360 can be thought of as completed automated since we apply auction optimizations even on the fixed CPM slice

Poirot



Advertiser surplus is a function of value - cost

	Impressions	Revenue
AdX	326,251,926	450,532
	332,397,428	467,690
	1.88%	3.81%
	[1.82, 1.95] %	[3.65, 3.97] %
3PE	185,614,226	625,828
	147,836,047	356,109
	-20.35%	-43.10%
	[-20.47, -20.23] %	[-43.21, -42.98] %


Boosts net revenue by \$150M

AdX 1p migration




Multi-calls that exploit the budget server

True Query Volume



Inflated Query Volume




*charts are for illustration only (not actual numbers)

- DBM is over 80% budget constrained
- Every query, a budget constrained ad participates in the auction with some budget throttling probability
- Exchanges send multiple calls for a single query to get multiple shots at budget throttling and land the largest bid
 - Money flows to publishers and exchanges that use repeated calls
 - Even if an exchange (say Adx or United) has the lowest margins, DBM might win through a different exchange (say Pubmatic) simply because the bid was 10x higher!

Project **Elmo** leverages cookie based budget throttling to ensure more consistent bids across calls

- We fix the advertisers that can purchase a query from a given cookie during any specific time_bucket (budget throttling based on cookie \times time_bucket)
- We bid the same across multiple calls since they occur close to each other in time
- Cookie usage goes deep (solution is somewhat easy in this case)



Supply path optimization

- The same inventory (e.g., nytimes.com) is available through several paths (AdX, Rubicon etc), and sometimes involving multiple hops (Rubicon buys, then re-sells to AdX)
- We have started receiving some data ([sellers.json](#) and [SupplyChain](#)) regarding the paths, though adoption of such standards are still quite low
- Regardless, SPO is a popular feature request from advertisers, and is an area of active investigation for us

Current priorities

1. Core optimization
 - a. Exchange bidding
 - b. Budget management
2. Lilypad/ Ribbit
 - a. Measurement
 - b. Fledge
 - c. Topics
3. Optimized targeting
4. Customizations
 - a. Custom bidding
 - b. Semantic matching

Google Marketing Platform

DV3 Optimized Targeting Review

Driving Performance with fully automated campaigns



September 01 2022
dbm-optimization

Summary

Proprietary + Confidential

- Optimized Targeting (OT) is a new targeting framework that empowers advertisers to leverage the power of Google ML for targeting to further improve their KPIs
- Developed with an eye on the cookieless world and a more efficient optimization framework, OT shifts the focus from binary **audience** eligibility decisions to bidding decisions across query and users
- The results show substantial improvements in advertiser performance across bidding strategies

[Learn more](#)

Problems with today's targeting technology

- Today's ad serving system built based on a manual-targeting-first mindset.
 - Targeting starts with "manual" targeting of groups of users (audiences) followed by expansions to larger and "similar" audiences
- This approach has many drawbacks:
 - The goal is to optimize advertiser KPI. User expansion is a prescriptive and suboptimal solution to this optimization problem
 - User expansions are made using several notions of user similarity, leading to dozens of complex models with problematic feedback loops in the system
 - The extent of expansion is ill-defined (slider controlled with no clear notion of optimal extent)
 - Post 3PCD, the notion of targeting a specific user group doesn't exist any more
- Today's auto-targeting technology has been in use in Google Ads for more than a decade
 - DV3 is the first Google buyer to employ this new auto-targeting technology.

Bird's eye view of Optimized Targeting

Determine optimal ad ranking &
auction bids to maximize advertiser KPI
Subject to Performance goal

Optimized targeting aims to solve the above optimization problem more directly (instead of being prescriptive about user expansions)


Advertisers observe the same KPI / dollar with optimized targeting as when targeting only using manual criteria

More Details

Google


Current Targeting Framework

- Advertisers provide us their audience input (e.g., in market for shoes, RMKT list).
- When opted-in, the Auto-Targeting/Similar-Audience system adds/removes Google Audience criteria (AT) or similar users (SA) on a daily/hourly basis based on whether they **perform** well.
- We still make 0-1 cutoff decisions for all advertisers once their “auto” criteria are added.
- Targeting Expansion systems are hand-tuned to trade-off revenue and computational resources.



Optimized Targeting

- Advertisers continue to provide us audience inputs (e.g., in market for shoes)
- The advertisers' ads will be considered eligible on all queries.
- The advertiser's audience input along with query and user interest-based signals are used to generate an ***optimal ad ranking and auction bid*** that maximizes advertiser KPI.



Optimized Targeting

- Why is OT a better technology?
 - Directly optimizes advertiser KPI without using the prescriptive “similar users/criteria” that are continuously added and removed
 - The solution is optimal, rather than a suboptimal 0-1 eligibility
 - System is significantly simpler and explainable
- How does OT manage computational cost?
 - Leans heavily on the Brasilia framework. Ultra-efficient ML models can handle large-scale ad candidate scorings.
 - [Project Phi](#) automatically optimizes the revenue-resource-latency trade-offs in the system
- What would the advertiser see?
 - Higher conversion (or other KPI) volume for the same spend compared to manual-only

Experiment Results

Google


Experiment Results: Buyside View

Prophetcore + Confidential

	Impressions	Revenue ^T
Total (Network)	0.5%	1.4%
DBM	2.6%	5.1%
GDA	-0.3%	-0.2%
UAC	-0.3%	-0.1%


^T Raw stats based on cookie-experiment stats @10% for 5 days.

[RASTA](#)



Rasta Reference:

https://experiments.corp.google.com/#/portal/experiments/analysis?label=_:uJkA7PW6nCeybh7CClwLpmf5CY



GA - AT: all LIs using Google Audience and auto targeting.

RMKT - SA: all LIs using RMKT and expansions (SA).

Mixed - AT/SA: all LIs using either RMKT or GA and enabled expansions.

Experiment Results: DV3 Performance Metrics			
Line Items opted-in targeting expansions			
	Spend	REMH CPD	Spend/2 + CPD
Max Conv (30% share of spend)	+40%	-4%	+16%
Max Clicks (14% share of spend)	+45%	+21%	+44%
Max ActiveView (8% share of spend)	+28%	+20%	+34%

T Active-View-PD is used for Max Active-View.

Google

https://docs.google.com/spreadsheets/d/1NGajdjFyukjVDxG0VIT-di2twzkOjgYs_Pp211bynA0/edit#gid=0